

Dominating Manipulations in Voting with Partial Information

Vincent Conitzer

Department of Computer Science
Duke University
Durham, NC 27708, USA
conitzer@cs.duke.edu

Toby Walsh

NICTA and UNSW
Sydney, Australia
toby.walsh@nicta.com.au

Lirong Xia

Department of Computer Science
Duke University
Durham, NC 27708, USA
lxia@cs.duke.edu

Abstract

We consider manipulation problems when the manipulator only has partial information about the votes of the non-manipulators. Such partial information is described by an *information set*, which is the set of profiles of the non-manipulators that are indistinguishable to the manipulator. Given such an information set, a *dominating manipulation* is a non-truthful vote that the manipulator can cast which makes the winner at least as preferable (and sometimes more preferable) as the winner when the manipulator votes truthfully. When the manipulator has full information, computing whether or not there exists a dominating manipulation is in P for many common voting rules (by known results). We show that when the manipulator has no information, there is no dominating manipulation for many common voting rules. When the manipulator's information is represented by partial orders and only a small portion of the preferences are unknown, computing a dominating manipulation is NP-hard for many common voting rules. Our results thus throw light on whether we can prevent strategic behavior by limiting information about the votes of other voters.

Introduction

In computational social choice, one appealing escape from the Gibbard-Satterthwaite theorem (Gibbard 1973; Satterthwaite 1975) was proposed in (Bartholdi, Tovey, and Trick 1989). Whilst manipulation may always be possible, perhaps it is computationally too difficult to find? Many results have subsequently been proven showing that various voting rules are NP-hard to manipulate (Bartholdi and Orlin 1991; Conitzer and Sandholm 2003; Elkind and Lipmaa 2005; Conitzer, Sandholm, and Lang 2007; Faliszewski, Hemaspaandra, and Schnoor 2008; Xia et al. 2009; Faliszewski, Hemaspaandra, and Schnoor 2010) in various senses. However, recent results suggest that computing a manipulation is easy on average or in many cases. Therefore, computational complexity seems to be a weak barrier against manipulation. See (Faliszewski, Hemaspaandra, and Hemaspaandra 2010; Faliszewski and Procaccia 2010) for some surveys of this recent research.

It is normally assumed that the manipulator has full information about the votes of the non-manipulators. The argument often given is that if it is NP-hard with full infor-

mation, then it only can be at least as computationally difficult with partial information. However, when there is only one manipulator, computing a manipulation is polynomial for most common voting rules, including all positional scoring rules, Copeland, maximin, and voting trees. The only known exceptions are STV (Bartholdi and Orlin 1991) and ranked pairs (Xia et al. 2009). Therefore, it is not clear whether a single manipulator has incentive to lie when the manipulator only has partial information.

In this paper, we study the problem of how one manipulator computes a manipulation based on partial information about the other votes. For example, the manipulator may know that some voters prefer one alternative to another, but might not be able to know all pairwise comparisons for all voters. We suppose the knowledge of the manipulator is described by an *information set E*. This is some subset of possible profiles of the non-manipulators which is known to contain the true profile. Given an information set and a pair of votes U and V , if for every profile in E , the manipulator is not worse off voting U than voting V , and there exists a profile in E such that the manipulator is strictly better off voting U , then we say that U *dominates* V . If there exists a vote U that dominates the true preferences of the manipulator then the manipulator has an incentive to vote untruthfully. We call this a *dominating manipulation*. If there is no such vote, then a risk-averse manipulator might have little incentive to vote strategically.

We are interested in whether a voting rule r is *immune* to dominating manipulations, meaning that a voter's true preferences are never dominated by another vote. If r is not immune to dominating manipulations, we are interested in whether r is *resistant*, meaning that computing whether a voter's true preferences are dominated by another vote U is NP-hard, or *vulnerable*, meaning that this problem is in P. These properties depend on both the voting rule and the form of the partial information. Interestingly, it is not hard to see that most voting rules are immune to manipulation when the partial information is just the current winner. For instance, with any majority consistent rule (for example, plurality), a risk averse manipulator will still want to vote for her most preferred alternative. This means that the chairman does not need to keep the current winner secret to prevent such manipulations. On the other hand, if the chairman lets slip more information, many rules stop being immune. With most scor-

ing rules, if the manipulator knows the current scores, then the rule is no longer immune to such manipulation. For instance, when her most preferred alternative is too far behind to win, the manipulator might vote instead for a less preferred candidate who can win.

In this paper, we focus on the case where the partial information is represented by a profile P_{po} of partial orders, and the information set E consists of all linear orders that extend P_{po} . The dominating manipulation problem is related to the *possible/necessary winner* problems (Konczak and Lang 2005; Walsh 2007; Betzler, Hemmann, and Niedermeier 2009; Betzler and Dorn 2010; Xia and Conitzer 2011). In possible/necessary winner problems, we are given an alternative c and a profile of partial orders P_{po} that represents the partial information of the voters' preferences. We are asked whether c is the winner for *some* extension of P_{po} (that is, c is a *possible winner*), or whether c is the winner for *every* extension of P_{po} (that is, c is a *necessary winner*). We note that in the possible/necessary winner problems, there is no manipulator and P_{po} represents the chair's partial information about the votes. In dominating manipulation problems, P_{po} represents the partial information of the manipulator about the non-manipulators.

We start with the special case where the manipulator has complete information. In this setting the dominating manipulation problem reduces to the standard manipulation problem, and many common voting rules are vulnerable to dominating manipulation (from known results). When the manipulator has no information, we show that a wide range of common voting rules are immune to dominating manipulation. When the manipulator's partial information is represented by partial orders, our results are summarized in Table 1.

DOMINATING MANIPULATION ¹	
STV	Resistant (Proposition 2)
Ranked pairs	Resistant (Proposition 2)
Borda	Resistant (Theorem 4)
Copeland	Resistant (Corollary 2)
Voting trees	Resistant (Corollary 2)
Maximin	Resistant (Theorem 7)
Plurality	Vulnerable (Algorithm 2)
Veto	Vulnerable (Omitted due to the space constraint.)

Table 1: Computational complexity of the dominating manipulation problems with partial orders, for common voting rules.

Our results are encouraging. For most voting rules r we study in this paper (except plurality and veto), hiding even a little information makes r resistant to dominating manipulation. If we hide all information, then r is immune to dominating manipulation. Therefore, limiting the information available to the manipulator appears to be a promising way to prevent strategic voting.

Preliminaries

Let $\mathcal{C} = \{c_1, \dots, c_m\}$ be the set of *alternatives* (or *candidates*). A linear order on \mathcal{C} is a transitive, antisymmetric,

¹All hardness results hold even when the number of undetermined pairs in each partial order is no more than a constant.

and total relation on \mathcal{C} . The set of all linear orders on \mathcal{C} is denoted by $L(\mathcal{C})$. An n -voter profile P on \mathcal{C} consists of n linear orders on \mathcal{C} . That is, $P = (V_1, \dots, V_n)$, where for every $j \leq n$, $V_j \in L(\mathcal{C})$. The set of all n -profiles is denoted by \mathcal{F}_n . We let m denote the number of alternatives. For any linear order $V \in L(\mathcal{C})$ and any $i \leq m$, $\text{Alt}(V, i)$ is the alternative that is ranked in the i th position in V . A *voting rule* r is a function that maps any profile on \mathcal{C} to a unique winning alternative, that is, $r : \mathcal{F}_1 \cup \mathcal{F}_2 \cup \dots \rightarrow \mathcal{C}$. The following are some common voting rules. In this paper, if not mentioned specifically, ties are broken in the fixed order $c_1 \succ c_2 \succ \dots \succ c_m$.

- **(Positional) scoring rules:** Given a *scoring vector* $\vec{s}_m = (\vec{s}_m(1), \dots, \vec{s}_m(m))$ of m integers, for any vote $V \in L(\mathcal{C})$ and any $c \in \mathcal{C}$, let $\vec{s}_m(V, c) = \vec{s}_m(j)$, where j is the rank of c in V . For any profile $P = (V_1, \dots, V_n)$, let $\vec{s}_m(P, c) = \sum_{j=1}^n \vec{s}_m(V_j, c)$. The rule will select $c \in \mathcal{C}$ so that $\vec{s}_m(P, c)$ is maximized. We assume scores are integers and decreasing. Some examples of positional scoring rules are *Borda*, for which the scoring vector is $(m-1, m-2, \dots, 0)$, *plurality*, for which the scoring vector is $(1, 0, \dots, 0)$, and *veto*, for which the scoring vector is $(1, \dots, 1, 0)$.

- **Copeland:** For any two alternatives c_i and c_j , we conduct a *pairwise election* in which we count how many voters rank c_i ahead of c_j , and how many rank c_j ahead of c_i . c_i wins if and only if the majority of voters rank c_i ahead of c_j . An alternative receives one point for each such win in a pairwise election. Typically, an alternative also receives half a point for each pairwise tie, but this will not matter for our results. The winner is the alternative with the highest score.

- **Maximin:** Let $D_P(c_i, c_j)$ be the number of votes that rank c_i ahead of c_j minus the number of votes that rank c_j ahead of c_i in the profile P . The winner is the alternative c that maximizes $\min\{D_P(c, c') : c' \in \mathcal{C}, c' \neq c\}$.

- **Ranked pairs:** This rule first creates an entire ranking of all the alternatives. In each step, we will consider a pair of alternatives c_i, c_j that we have not previously considered; specifically, we choose the remaining pair with the highest $D_P(c_i, c_j)$. We then fix the order $c_i \succ c_j$, unless this contradicts previous orders that we fixed (that is, it violates transitivity). We continue until we have considered all pairs of alternatives (hence we have a full ranking). The alternative at the top of the ranking wins.

- **Voting trees:** A voting tree is a binary tree with m leaves, where each leaf is associated with an alternative. In each round, there is a pairwise election between an alternative c_i and its sibling c_j : if the majority of voters prefer c_i to c_j , then c_j is eliminated, and c_i is associated with the parent of these two nodes. The alternative that is associated with the root of the tree (i.e. wins all its rounds) is the winner.

- **Single transferable vote (STV):** The election has m rounds. In each round, the alternative that gets the lowest plurality score (the number of times that the alternative is ranked in the top position) drops out, and is removed from all of the votes (so that votes for this alternative transfer to another alternative in the next round). The last-remaining alternative is the winner.

For any profile P , we let $\text{WMG}(P)$ denote the *weighted majority graph* of P , defined as follows. $\text{WMG}(P)$ is a di-

rected graph whose vertices are the alternatives. For $i \neq j$, if $D_P(c_i, c_j) > 0$, then there is an edge (c_i, c_j) with weight $w_{ij} = D_P(c_i, c_j)$.

We say that a voting rule r is based on the *weighted majority graph* (WMG), if for any pair of profiles P_1, P_2 such that $\text{WMG}(P_1) = \text{WMG}(P_2)$, we have $r(P_1) = r(P_2)$. A voting rule r is *Condorcet consistent* if it always selects the Condorcet winner (that is, the alternative that wins each of its pairwise elections) whenever one exists.

Manipulation with Partial Information

We now introduce the framework of this paper. Suppose there are $n \geq 1$ non-manipulators and one manipulator. The information the manipulator has about the votes of the non-manipulators is represented by an *information set* E . The manipulator knows for sure that the profile of the non-manipulators is in E . However, the manipulator does not know exactly which profile in E it is. Usually E is represented in a compact way. Let \mathcal{I} denote the set of all possible information sets in which the manipulator may find herself.

Example 1. Suppose the voting rule is r .

- If the manipulator has no information, then the only information set is $E = \mathcal{F}_n$. Therefore $\mathcal{I} = \{\mathcal{F}_n\}$.
- If the manipulator has complete information, then $\mathcal{I} = \{\{P\} : P \in \mathcal{F}_n\}$.
- If the manipulator knows the current winner (before the manipulator votes), then the set of all information sets the manipulator might know is $\mathcal{I} = \{E_1, E_2, \dots, E_m\}$, where for any $i \leq m$, $E_i = \{P \in \mathcal{F}_n : r(P) = c_i\}$.

Let V_M denote the true preferences of the manipulator. Given a voting rule r and an information set E , we say that a vote U *dominates* another vote V , if for every profile $P \in E$, we have $r(P \cup \{U\}) \succeq_{V_M} r(P \cup \{V\})$, and there exists $P' \in E$ such that $r(P' \cup \{U\}) \succ_{V_M} r(P' \cup \{V\})$. In other words, when the manipulator only knows the voting rule r and the fact that the profile of the non-manipulators is in E (and no other information), voting U is a strategy that dominates voting V . We define the following two decision problems.

Definition 1. Given a voting rule r , an information set E , the true preferences V_M of the manipulator, and two votes V and U , we are asked the following two questions.

- Does U dominate V ? This is the DOMINATION problem.
- Does there exist a vote V' that dominates V_M ? This is the DOMINATING MANIPULATION problem.

We stress that usually E is represented in a compact way, otherwise the input size would already be exponentially large, which would trivialize the computational problems. Given a set \mathcal{I} of information sets, we say a voting rule r is *immune* to dominating manipulation, if for every $E \in \mathcal{I}$ and every V_M that represents the manipulator's preferences, V_M is not dominated; r is *resistant* to dominating manipulation, if DOMINATING MANIPULATION is NP-hard (which means that r is not immune to dominating manipulation, assuming P ≠ NP); and r is *vulnerable* to dominating manipulation, if r is not immune to dominating manipulation, and DOMINATING MANIPULATION is in P.

Manipulation with Complete/No Information

In this section we focus on the following two special cases: (1) the manipulator has complete information, and (2) the manipulator has no information. It is not hard to see that when the manipulator has complete information, DOMINATING MANIPULATION coincides with the standard manipulation problem. Therefore, our framework of dominating manipulation is an extension of the traditional manipulation problem, and we immediately obtain the following proposition from the Gibbard-Satterthwaite theorem (Gibbard 1973; Satterthwaite 1975).

Proposition 1. When $m \geq 3$ and the manipulator has full information, a voting rule satisfies non-imposition and is immune to dominating manipulation if and only if it is a dictatorship.

The following proposition directly follows from the computational complexity of the manipulation problems for some common voting rules (Bartholdi, Tovey, and Trick 1989; Bartholdi and Orlin 1991; Conitzer, Sandholm, and Lang 2007; Zuckerman, Procaccia, and Rosenschein 2009; Xia et al. 2009).

Proposition 2. When the manipulator has complete information, STV and ranked pairs are resistant to DOMINATING MANIPULATION; all positional scoring rules, Copeland, voting trees, and maximin are vulnerable to dominating manipulation.

Next, we investigate the case where the manipulator has no information. We obtain the following positive results. Due to the space constraint, most proofs are omitted.

Theorem 1. When the manipulator has no information, any Condorcet consistent voting rule r is immune to dominating manipulation.

Theorem 2. When the manipulator has no information, Borda is immune to dominating manipulation.

Theorem 3. When the manipulator has no information and $n \geq 6(m - 2)$, any positional scoring rule is immune to dominating manipulation.

These results demonstrate that the information that the manipulator has about the votes of the non-manipulators plays an important role in determining strategic behavior. When the manipulator has complete information, many common voting rules are vulnerable to dominating manipulation, but if the manipulator has no information, then many common voting rules become immune to dominating manipulation.

Manipulation with Partial Orders

In this section, we study the case where the manipulator has partial information about the votes of the non-manipulators. We suppose the information is represented by a profile P_{po} composed of partial orders. That is, the information set is $E = \{P \in \mathcal{F}_n : P \text{ extends } P_{po}\}$. We note that the two cases discussed in the previous section (complete information and no information) are special cases of manipulation with partial orders. Consequently, by Proposition 1, when the manipulator's information is represented by partial orders and $m \geq 3$, no voting rule that satisfies non-imposition and non-dictatorship is immune to dominating manipulation. It also

follows from Theorem 2 that STV and ranked pairs are resistant to dominating manipulation. The next theorem states that even when the manipulator only misses a tiny portion of the information, Borda becomes resistant to dominating manipulation.

Theorem 4. DOMINATION and DOMINATING MANIPULATION with partial orders are NP-hard for Borda, even when the number of unknown pairs in each vote is no more than 4.

Proof. We only prove that DOMINATION is NP-hard, via a reduction from EXACT COVER BY 3-SETS (X3C). The proof for DOMINATING MANIPULATION is omitted due to space constraint. The reduction is similar to the proof of the NP-hardness of the possible winner problems under positional scoring rules in (Xia and Conitzer 2011).

In an X3C instance, we are given two sets $\mathcal{V} = \{v_1, \dots, v_q\}$, $\mathcal{S} = \{S_1, \dots, S_t\}$, where for any $j \leq t$, $S_j \subseteq \mathcal{V}$ and $|S_j| = 3$. We are asked whether there exists a subset \mathcal{S}' of \mathcal{S} such that each element in \mathcal{V} is in exactly one of the 3-sets in \mathcal{S}' . We construct a DOMINATION instance as follows.

Alternatives: $\mathcal{C} = \{c, w, d\} \cup \mathcal{V}$, where d is an auxiliary alternative. Therefore, $m = |\mathcal{C}| = q + 3$. Ties are broken in the following order: $c \succ w \succ \mathcal{V} \succ d$.

Manipulator's preferences and possible manipulation: $V_M = [w \succ c \succ d \succ \mathcal{V}]$. We are asked whether $V = V_M$ is dominated by $U = [w \succ d \succ c \succ \mathcal{V}]$.

The profile of partial orders: Let $P_{po} = P_1 \cup P_2$, defined as follows.

First part (P_1) of the profile: For each $j \leq t$, We define a partial order O_j as follows.

$$O_j = [w \succ S_j \succ d \succ \text{Others}] \setminus [\{w\} \times (S_j \cup \{d\})]$$

That is, O_j is a partial order that agrees with $w \succ S_j \succ d \succ \text{Others}$, except that the pairwise relations between (w, S_j) and (w, d) are not determined (and these are the only 4 unknown relations). Let $P_1 = \{O_1, \dots, O_t\}$.

Second part (P_2) of the profile: We first give the properties that we need P_2 to satisfy, then show how to construct P_2 in polynomial time. All votes in P_2 are linear orders that are used to adjust the score differences between alternatives. Let $P'_1 = \{w \succ S_i \succ d \succ \text{Others} : i \leq t\}$. That is, P'_1 ($|P'_1| = t$) is an extension of P_1 (in fact, P'_1 is the set of linear orders that we started with to obtain P_1 , before removing some of the pairwise relations). Let $\vec{s}_m = (m - 1, \dots, 0)$. P_2 is a set of linear orders such that the following holds for $Q = P'_1 \cup P_2 \cup \{V\}$:

(1) For any $i \leq q$, $\vec{s}_m(Q, c) - \vec{s}_m(Q, v_i) = 1$, $\vec{s}_m(Q, w) - \vec{s}_m(Q, c) = 4q/3$.

(2) For any $i \leq q$, the scores of v_i and w, c are higher than the score of d in any extension of $P_1 \cup P_2 \cup \{V\}$ and in any extension of $P_1 \cup P_2 \cup \{U\}$.

(3) The size of P_2 is polynomial in $t + q$.

We now show how to construct P_2 in polynomial time. For any alternative $a \neq d$, we define the following two votes: $W_a = \{[a \succ d \succ \text{Others}], [\text{Rev(Others)} \succ a \succ d]\}$, where Rev(Others) is the reversed order of the alternatives in $\mathcal{C} \setminus \{a, d\}$. We note that for any alternative $a' \in \mathcal{C} \setminus \{a, d\}$, $\vec{s}_m(W, a) - \vec{s}_m(W, a') = 1$ and $\vec{s}_m(W, a') - \vec{s}_m(W, d) = 1$. Let $Q_1 = P'_1 \cup \{V\}$. P_2 is composed of the following parts:

(1) $tm - \vec{s}_m(Q_1, c)$ copies of W_c .

(2) $tm + 4q/3 - \vec{s}_m(Q_1, w)$ copies of W_w .

(2) For each $i \leq q$, there are $tm - 1 - \vec{s}_m(Q_1, v_i)$ copies of W_{v_i} .

We next prove that V is dominated by U if and only if c is the winner in at least one extension of $P_{po} \cup \{V\}$. We note that for any $v \in \mathcal{V} \cup \{w\}$, the score of v in V is the same as the score of v in U . The score of c in U is lower than the score of c in V . Therefore, for any extension P^* of P_{po} , if $r(P^* \cup \{V\}) \in (\{w\} \cup \mathcal{V})$, then $r(P^* \cup \{V\}) = r(P^* \cup \{U\})$ (because d cannot win). Hence, for any extension P^* of P_{po} , voting U can result in a different outcome than voting V only if $r(P^* \cup V) = c$. If there exists an extension P^* of P_{po} such that $r(P^* \cup \{V\}) = c$, then we claim that the manipulator is strictly better off voting U than voting V . Let P_1^* denote the extension of P_1 in P^* . Then, because the total score of w is no more than the total score of c , w is ranked lower than d at least $\frac{q}{3}$ times in P_1^* . Meanwhile, for each $i \leq q$, v_i is not ranked higher than w more than one time in P_1^* , because otherwise the total score of v_i will be strictly higher than the total score of c . That is, the votes in P_1^* where $d \succ w$ make up a solution to the X3C instance. Therefore, the only possibility for c to win is for the scores of c, w , and all alternatives in \mathcal{V} to be the same (so that c wins according to the tie-breaking mechanism). Now, we have $w = r(P^* \cup \{U\})$. Because $w \succ_{V_M} c$, the manipulator is better off voting U . It follows that V is dominated by U if and only if there exists an extension of $P_{po} \cup \{V\}$ where c is the winner.

The above reasoning also shows that V is dominated by U if and only if the X3C instance has a solution. Therefore, DOMINATION is NP-hard. \square

Theorem 4 can be generalized to a class of scoring rules similar to the class of rules in Theorem 1 in (Xia and Conitzer 2011), which does not include plurality or veto. In fact, as we will show later, plurality and veto are vulnerable to dominating manipulation.

We now investigate the relationship to the possible winner problem in more depth. In a possible winner problem (r, P_{po}, c) , we are given a voting rule r , a profile P_{po} composed of n partial orders, and an alternative c . We are asked whether there exists an extension P of P_{po} such that $c = r(P)$. Intuitively, both DOMINATION and DOMINATING MANIPULATION seem to be harder than the possible winner problem under the same rule. Next, we present two theorems, which show that for any WMG-based rule, DOMINATION and DOMINATING MANIPULATION are harder than two special possible winner problems, respectively.

We first define a notion that will be used in defining the two special possible winner problems. For any instance of the possible winner problem (r, P_{po}, c) , we define its WMG partition $\mathcal{R} = \{R_{c'} : c' \in \mathcal{C}\}$ as follows. For any $c' \in \mathcal{C}$, let $R_{c'} = \{\text{WMG}(P) : P \text{ extends } P_{po} \text{ and } r(P) = c'\}$. That is, $R_{c'}$ is composed of all WMGs of the extensions of P_{po} , where the winner is c' . It is possible that for some $c' \in \mathcal{C}$, $R_{c'}$ is empty. For any subset $\hat{\mathcal{C}}' \subseteq \mathcal{C} \setminus \{c\}$, we let $G_{\hat{\mathcal{C}}'}$ denote the weighted majority graph where for each $c' \in \hat{\mathcal{C}}'$, there is an edge $c' \rightarrow c$ with weight 2, and these are the only edges in $G_{\hat{\mathcal{C}}'}$. We are ready to define the two special possible winner problems for WMG-based voting rules.

Definition 2. Let d^* be an alternative and let \mathcal{C}' be a

nonempty subset of $\mathcal{C} \setminus \{c, d^*\}$. For any WMG-based voting rule r , we let $PW_1(d^*, \mathcal{C}')$ denote the set of possible winner problems (r, P_{po}, c) satisfying the following conditions:

1. For any $G \in R_c$, $r(G + G_{\mathcal{C}'}) = d^*$.
2. For any $c' \neq c$ and any $G \in R_{c'}$, $r(G + G_{\mathcal{C}'}) = r(G)$.
3. For any $c' \in \mathcal{C}'$, $R_{c'} = \emptyset$.

We recall that R_c and $R_{c'}$ are elements in the WMG partition of the possible winner problem.

Definition 3. Let d^* be an alternative and let \mathcal{C}' be a nonempty subset of $\mathcal{C} \setminus \{c, d^*\}$. For any WMG-based voting rule r , we let $PW_2(d^*, \mathcal{C}')$ denote the problem instances (r, P_{po}, c) of $PW_1(d^*, \mathcal{C}')$, where for any $c' \in \mathcal{C} \setminus \{c, d^*\}$, $R_{c'} = \emptyset$.

Theorem 5. Let r be a WMG-based voting rule. There is a polynomial time reduction from $PW_1(d^*, \mathcal{C}')$ to DOMINATION with partial orders, both under r .

Proof. Let (r, P_{po}, c) be a $PW_1(d^*, \mathcal{C}')$ instance. We construct the following DOMINATION instance. Let the profile of partial orders be $Q_{po} = P_{po} \cup \{\text{Rev}(d^* \succ c \succ \mathcal{C}' \succ \text{Others})\}$, $V = V_M = [d^* \succ c \succ \mathcal{C}' \succ \text{Others}]$, and $U = [d^* \succ \mathcal{C}' \succ c \succ \text{Others}]$. Let P be an extension of P_{po} . It follows that $\text{WMG}(P \cup \{\text{Rev}(d^* \succ c \succ \mathcal{C}' \succ \text{Others}), V\}) = \text{WMG}(P)$, and $\text{WMG}(P \cup \{\text{Rev}(d^* \succ c \succ \mathcal{C}' \succ \text{Others}), U\}) = \text{WMG}(P) + G_{\mathcal{C}'}$. Therefore, the manipulator can change the winner if and only if $\text{WMG}(P) \in R_c$, which is equivalent to c being a possible winner. We recall that by the definition of $PW_1(d^*, \mathcal{C}')$, for any $G \in R_c$, $r(G + G_{\mathcal{C}'}) = d^*$; for any $c' \neq c$ and any $G \in R_{c'}$, $r(G + G_{\mathcal{C}'}) = c'$; and $d^* \succ_V c$. It follows that $V (=V_M)$ is dominated by U if and only if the $PW_1(d^*, \mathcal{C}')$ instance has a solution. \square

Theorem 5 can be used to prove that DOMINATION is NP-hard for Copeland, maximin, and voting trees, even when the number of undetermined pairs in each partial order is bounded above by a constant. It suffices to show that for each of these rules, there exist d^* and \mathcal{C}' such that $PW_1(d^*, \mathcal{C}')$ is NP-hard. To prove this, we can modify the NP-completeness proofs of the possible winner problems for Copeland, maximin, and voting trees by Xia and Conitzer (Xia and Conitzer 2011). These proofs are omitted due to space constraint.

Corollary 1. DOMINATION with partial orders is NP-hard for Copeland, maximin, and voting trees, even when the number of unknown pairs in each vote is bounded above by a constant.

Theorem 6. Let r be a WMG-based voting rule. There is a polynomial-time reduction from $PW_2(d^*, \mathcal{C}')$ to DOMINATING MANIPULATION with partial orders, both under r .

Proof. The proof is similar to the proof for Theorem 5. We note that d^* is the manipulator's top-ranked alternative. Therefore, if c is not a possible winner, then $V (=V_M)$ is not dominated by any other vote; if c is a possible winner, then V is dominated by $U = [w \succ \mathcal{C}' \succ c \succ \text{Others}]$. \square

Similarly, we have the following corollary.

Corollary 2. DOMINATING MANIPULATION with partial orders is NP-hard for Copeland and voting trees, even when the number of unknown pairs in each vote is bounded above by a constant.

It is an open question if $PW_2(d^*, \mathcal{C}')$ with partial orders is NP-hard for maximin. However, we can directly prove that DOMINATING MANIPULATION is NP-hard for maximin by a reduction from x3C.

Theorem 7. DOMINATING MANIPULATION with partial orders is NP-hard for maximin, even when the number of unknown pairs in each vote is no more than 4.

For plurality and veto, there exist polynomial-time algorithms for both DOMINATION and DOMINATING MANIPULATION. Given an instance of DOMINATION, denoted by (r, P_{po}, V_M, V, U) , we say that U is a possible improvement of V , if there exists an extension P of P_{po} such that $r(P \cup \{U\}) \succ_{V_M} r(P \cup \{V\})$. It follows that U dominates V if and only if U is a possible improvement of V , and V is not a possible improvement of U . We first introduce an algorithm (Algorithm 1) that checks whether U is a possible improvement of V for plurality.

Let c_{i^*} (resp., c_{j^*}) denote the top-ranked alternative in V (resp., U). We will check whether there exists $0 \leq l \leq n$, $d, d' \in \mathcal{C}$ with $d' \succ_{V_M} d$, and an extension P^* of P_{po} , such that if the manipulator votes for V , then the winner is d , whose plurality score in P^* is l , and if the manipulator votes for U , then the winner is d' . We note that if such d, d' exist, then either $d = c_{i^*}$ or $d' = c_{j^*}$ (or both hold). To this end, we solve multiple maximum-flow problems defined as follows.

Let $\mathcal{C}' \subset \mathcal{C}$ denote a set of alternatives. Let $\vec{e} = (e_1, \dots, e_m) \in \mathbb{N}^m$ be an arbitrary vector composed of m natural numbers such that $\sum_{i=1}^m e_i \geq n$. We define a maximum-flow problem $F_{\mathcal{C}'}^{\vec{e}}$ as follows.

Vertices: $\{s, O_1, \dots, O_n, c_1, \dots, c_m, y, t\}$.

Edges:

- For any O_i , there is an edge from s to O_i with capacity 1.
- For any O_i and c_j , there is an edge $O_i \rightarrow c_j$ with capacity 1 if and only if c_j can be ranked in the top position in at least one extension of O_i .
- For any $c_i \in \mathcal{C}'$, there is an edge $c_i \rightarrow t$ with capacity e_i .
- For any $c_i \in \mathcal{C} \setminus \mathcal{C}'$, there is an edge $c_i \rightarrow y$ with capacity e_i .
- There is an edge $y \rightarrow t$ with capacity $n - \sum_{c_i \in \mathcal{C}'} e_i$.

For example, $F_{\{c_1, c_2\}}^{\vec{e}}$ is illustrated in Figure 1.




Figure 1: $F_{\{c_1, c_2\}}^{\vec{e}}$.

It is not hard to see that $F_{\mathcal{C}'}^{\vec{e}}$ has a solution whose value is n if and only if there exists an extension P^* of P_{po} , such that (1) for each $c_i \in \mathcal{C}'$, the plurality of c_i is exactly e_i , and

(2) for each $c_{i'} \notin \mathcal{C}'$, the plurality of $c_{i'}$ is no more than $e_{i'}$. Now, for any pair of alternatives $d = c_i, d' = c_j$ such that $d' \succ_{V_M} d$ and either $d = c_{i^*}$ or $d' = c_{j^*}$, we define the set of *admissible maximum-flow problems* A_{Plu}^l to be the set of maximum flow problems $F_{c_i, c_j}^{\vec{e}}$ where $e_i = l$, and if $F_{c_i, c_j}^{\vec{e}}$ has a solution, then the manipulator can improve the winner by voting for U . Details are omitted due to space constraint. Algorithm 1 solves all maximum-flow problems in A_{Plu}^l to check whether U is a possible improvement of V .

Algorithm 1: PossibleImprovement(V, U)

```

1 Let  $c_{i^*} = \text{Alt}(V, 1)$  and  $c_{j^*} = \text{Alt}(U, 1)$ .
2 for any  $0 \leq l \leq n$  and any pair of alternatives
    $d = c_i, d' = c_j$  such that  $d' \succ_{V_M} d$  and either  $d = c_{i^*}$ 
   or  $d' = c_{j^*}$  do
3   Compute  $A_{\text{Plu}}^l$ .
4   for each maximum-flow problem  $F_{\mathcal{C}'}^{\vec{e}}$  in  $A_{\text{Plu}}^l$  do
5     if  $\sum_{c_i \in \mathcal{C}'} e_i \leq n$  and the value of maximum
        flow in  $F_{\mathcal{C}'}^{\vec{e}}$  is  $n$  then
6       Output that the  $U$  is a possible improvement
          of  $V$ , terminate the algorithm.
7     end
8   end
9 end
10 Output that  $U$  is not a possible improvement of  $V$ .

```

The algorithm for DOMINATION (Algorithm 2) runs Algorithm 1 twice to check whether U is a possible improvement of V , and whether V is a possible improvement of U . The algorithm for DOMINATING MANIPULATION for

Algorithm 2: Domination

```

1 if PossibleImprovement( $V, U$ )= “yes” and
   PossibleImprovement( $U, V$ )= “no” then
2   | Output that  $V$  is dominated by  $U$ .
3 end
4 else
5   | Output that  $V$  is not dominated by  $U$ .
6 end

```

plurality simply runs Algorithm 2 $m - 1$ times. In the input we always have that $V = V_M$, and for each alternative in $\mathcal{C} \setminus \{\text{Alt}(V, 1)\}$, we solve an instance where that alternative is ranked first in U . If in any step V is dominated by U , then there is a dominating manipulation; otherwise V is not dominated by any other vote. The algorithms for DOMINATION and DOMINATING MANIPULATION for veto are similar. We omit the details due to space constraint.

Future Work

Analysis of manipulation with partial information provides insight into what needs to be kept confidential in an election. For instance, in a plurality or veto election, revealing (perhaps unintentionally) part of the preferences of non-manipulators may open the door to strategic voting. An interesting open question is whether there are any more general relationships between the possible winner problem and the dominating manipulation problem with partial orders. It

would be interesting to identify cases where voting rules are resistant or even immune to manipulation based on other types of partial information, for example, the set of possible winners. We may also consider other types of strategic behavior with partial information in our framework, for example, coalitional manipulation, bribery, and control. We are currently working on proving completeness results for higher levels of the polynomial hierarchy for problems similar to those studied in this paper.

Acknowledgments

Vincent Conitzer and Lirong Xia acknowledge NSF CAREER 0953756 and IIS-0812113, and an Alfred P. Sloan fellowship for support. Toby Walsh is supported by the Australian Department of Broadband, Communications and the Digital Economy, the ARC, and the Asian Office of Aerospace Research and Development (AOARD-104123). Lirong Xia is supported by a James B. Duke Fellowship. We thank all AAAI-11 reviewers for their helpful comments and suggestions.

References

- [Bartholdi and Orlin 1991] Bartholdi, III, J., and Orlin, J. 1991. Single transferable vote resists strategic voting. *SCW* 8(4):341–354.
- [Bartholdi, Tovey, and Trick 1989] Bartholdi, III, J.; Tovey, C.; and Trick, M. 1989. The computational difficulty of manipulating an election. *SCW* 6(3):227–241.
- [Betzler and Dorn 2010] Betzler, N., and Dorn, B. 2010. Towards a dichotomy for the possible winner problem in elections based on scoring rules. *JCSS* 76(8):812–836.
- [Betzler, Hemmann, and Niedermeier 2009] Betzler, N.; Hemmann, S.; and Niedermeier, R. 2009. A multivariate complexity analysis of determining possible winners given incomplete votes. In *Proc. IJCAI*, 53–58.
- [Conitzer and Sandholm 2003] Conitzer, V., and Sandholm, T. 2003. Universal voting protocol tweaks to make manipulation hard. In *Proc. IJCAI*, 781–788.
- [Conitzer, Sandholm, and Lang 2007] Conitzer, V.; Sandholm, T.; and Lang, J. 2007. When are elections with few candidates hard to manipulate? *JACM* 54(3):1–33.
- [Elkind and Lipmaa 2005] Elkind, E., and Lipmaa, H. 2005. Hybrid voting protocols and hardness of manipulation. In *Proc. ISAAC*, 24–26.
- [Faliszewski and Procaccia 2010] Faliszewski, P., and Procaccia, A. D. 2010. AI’s war on manipulation: Are we winning? *AI Magazine* 31(4):53–64.
- [Faliszewski, Hemaspaandra, and Hemaspaandra 2010] Faliszewski, P.; Hemaspaandra, E.; and Hemaspaandra, L. A. 2010. Using complexity to protect elections. *Commun. ACM* 53:74–82.
- [Faliszewski, Hemaspaandra, and Schnoor 2008] Faliszewski, P.; Hemaspaandra, E.; and Schnoor, H. 2008. Copeland voting: Ties matter. In *Proc. AAMAS*, 983–990.
- [Faliszewski, Hemaspaandra, and Schnoor 2010] Faliszewski, P.; Hemaspaandra, E.; and Schnoor, H. 2010. Manipulation of copeland elections. In *Proc. AAMAS*, 367–374.
- [Gibbard 1973] Gibbard, A. 1973. Manipulation of voting schemes: A general result. *Econometrica* 41:587–602.
- [Konczak and Lang 2005] Konczak, K., and Lang, J. 2005. Voting procedures with incomplete preferences. In *Multidisciplinary Workshop on Advances in Preference Handling*.
- [Satterthwaite 1975] Satterthwaite, M. 1975. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems

for voting procedures and social welfare functions. *JET* 10:187–217.

[Walsh 2007] Walsh, T. 2007. Uncertainty in preference elicitation and aggregation. In *Proc. AAAI*, 3–8.

[Xia and Conitzer 2011] Xia, L., and Conitzer, V. 2011. Determining possible and necessary winners under common voting rules given partial orders. To appear in *JAIR*.

[Xia et al. 2009] Xia, L.; Zuckerman, M.; Procaccia, A. D.; Conitzer, V.; and Rosenschein, J. 2009. Complexity of unweighted coalitional manipulation under some common voting rules. In *Proc. IJCAI*, 348–353.

[Zuckerman, Procaccia, and Rosenschein 2009] Zuckerman, M.; Procaccia, A. D.; and Rosenschein, J. S. 2009. Algorithms for the coalitional manipulation problem. *Artificial Intelligence* 173(2):392–412.